

Convective regularization for optical flow

José A. Iglesias*, Clemens Kirisits†

October 14, 2015

Abstract

We argue that the time derivative in a fixed coordinate frame may not be the most appropriate measure of time regularity of an optical flow field. Instead, for a given velocity field v we consider the convective acceleration $v_t + \nabla v v$ which describes the acceleration of objects moving according to v . Consequently we investigate the suitability of the nonconvex functional $\|v_t + \nabla v v\|_{L^2}^2$ as a regularization term for optical flow. We demonstrate that this term acts as both a spatial and a temporal regularizer and has an intrinsic edge-preserving property. We incorporate it into a contrast invariant and time-regularized variant of the Horn-Schunck functional, prove existence of minimizers and verify experimentally that it addresses some of the problems of basic quadratic models. For the minimization we use an iterative scheme that approximates the original nonlinear problem with a sequence of linear ones. We believe that the convective acceleration may be gainfully introduced in a variety of optical flow models.

1 Introduction

Motivation. Optical flow is the apparent motion in a sequence of images and can be described by a velocity field. Using variational techniques to estimate this velocity field requires the design of an appropriate energy functional. Typically such a functional is a sum of two parts. The first part, sometimes called data term, measures the accuracy with which the velocity field describes the observable image motion. By ensuring consistency of the flow field, the second part, also called regularization term, gives information the data term cannot provide and thereby addresses the inherent ill-posedness of the optical flow problem. Naturally, related research has to a large degree been concerned with tuning the second term to optimally capture the characteristics of velocities of real-world image sequences.

Generally speaking regularization terms for optical flow fall into two categories: those which only penalize spatial derivatives of the velocity field, and those where derivatives in both space and time are penalized. The first category is by far the more popular one. Its developments have closely followed those of variational image restoration, because the regularity of images dictates to a significant degree the spatial regularity of velocity fields describing their motion.

*Computational Science Center, University of Vienna, Vienna, Austria

†Johann Radon Institute for Computational and Applied Mathematics (RICAM), Austrian Academy of Sciences, Linz, Austria

In other words, image discontinuities tend to coincide with motion discontinuities. These, it was found, can be appropriately dealt with using subquadratic or anisotropic regularization.

Research on the second category, i.e. that of spatiotemporal regularization, is relatively scarce. Again this is in analogy with the field of image processing, where the joint restoration of image sequences is quite under-represented as compared to the restoration of single frames. The inclusion of time derivatives has the obvious disadvantage of turning a series of decoupled 2D problems into one 3D problem. In the early days of computer vision limited computer memory and processing power was prohibitive to such approaches. Clearly this is no longer the case. On the contrary, recent publications have articulated and addressed the need for time-coupled models even for data with three instead of two space dimensions. See for example [2, 15].

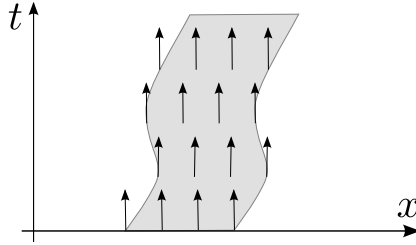


Figure 1: Using partial time derivatives in flow regularization can blur across object boundaries. Here, the shaded area represents the space-time signature of a moving object and the time direction points both inside and outside of it at different points.

It is well-known that indiscriminate smoothing in space can lead to blurring at the boundaries of objects. This statement is equally true in the space-time domain, if we identify objects with their space-time signature (Figure 1). In particular, blind penalization of the partial time derivative of a velocity field can lead to a loss of accuracy, especially at the boundaries of *moving* objects. Therefore we suggest to take the time derivative only along the movement of the object.

This derivative has a very natural physical interpretation. Let $v(t, x)$ be a time dependent velocity field on \mathbb{R}^d , and let $t \mapsto \gamma(t; x_0)$ be the trajectory of a certain object initially located at $x_0 \in \mathbb{R}^d$ that moves according to v . The translation into mathematical terms of this connection between γ and v is the following initial value problem

$$\begin{aligned}\gamma'(t; x_0) &= v(t, \gamma(t; x_0)) \\ \gamma(0; x_0) &= x_0.\end{aligned}\tag{1}$$

Taking time derivatives on both sides of the differential equation above gives a formula in terms of v for the acceleration of the object moving along γ

$$\begin{aligned}\gamma''(t; x_0) &= \frac{d}{dt}v(t, \gamma(t; x_0)) \\ &= \frac{\partial}{\partial t}v(t, \gamma(t; x_0)) + \nabla v(t, \gamma(t; x_0))v(t, \gamma(t; x_0))\end{aligned}$$

This expression, which is sometimes called *convective acceleration*, is to be contrasted with the partial time derivative of v evaluated at $(t, \gamma(t; x_0))$

$$\frac{\partial}{\partial t} v(t, \gamma(t; x_0)),$$

which has a notably different physical meaning. In this article we argue that in some situations the former can be more appropriate for the regularization of optical flow.

This work is devoted to investigating the suitability of the functional

$$v \mapsto \|v_t + \nabla v v\|_{L^2}^2 \quad (2)$$

as a smoothness term for optical flow. We believe a sensible way to do so is to include this term into a simple optical flow model and to compare the results with those of the same model without the new term. The reference model we choose for this task is a contrast invariant and time-regularized version of the Horn-Schunck functional [7]. Incorporating (2) into our model comes, however, at the cost of leading to a nonconvex functional and to nonlinear optimality conditions. Numerically, we approximate the problem by a sequence of convex quadratic ones.

We stress that it is not the aim of the proposed model to compete in the highest ranks of motion estimation benchmarks. Instead we want to point out a certain aspect of optical flow regularization, which we feel has not received its due amount of attention. Even though it leads to a more challenging variational problem, we demonstrate that a reasonable approximation is numerically tractable and pays off in terms of accuracy.

Related work. Horn and Schunck [7] are generally credited with having laid the groundwork for variational optical flow models. Penalizing the squared L^2 -norms of the spatial derivatives, their model regularizes the optical flow isotropically as well as homogeneously over the image domain. In [10] Nagel proposed to suppress smoothing across image discontinuities by only penalizing the derivatives of the vector field along the level lines of the image. This idea was subsequently extended to the space-time domain in [11].

A different route to spatiotemporal regularization was taken by Weickert and Schnörr. In [19] they used an isotropic but essentially subquadratic regularizer. While still convex this approach leads to nonlinear optimality conditions. In the survey [18] the same authors classify convex spatial as well as spatiotemporal regularizers for optical flow. According to their nomenclature our approach would classify as anisotropic and flow-driven. It is, however, fundamentally different from the regularizers considered there, not only because of nonconvexity. In [18] the anisotropy of flow-driven regularizers is determined by the Jacobian of the velocity field, whereas in our approach it is determined by the velocity itself.

Chaudhury and Mehrotra pursue an interesting trajectory-based regularization strategy [3]. Motivated by the principle of least action and the inertia of motion they postulate that both length and curvature of trajectories of moving objects should be minimal. There is a close connection between a trajectory's curvature and the convective acceleration of an object moving along that trajectory. We discuss this connection in the next section. More recently, Salgado

and Sánchez [14] as well as Volz and coauthors [17] have proposed time-discrete models using explicit trajectorial regularization. The present article tries to capture the essence of trajectorial regularization for optical flow in an entirely continuous setting.

Finally, we note that for the problem of image sequence reconstruction it is not uncommon to regularize the sequence along the optical flow, which can be either precomputed or estimated simultaneously. This derivative along the optical flow is just the convective derivative of the image sequence. See [4, 12] for example.

Outline. In Section 2 we introduce our model. In Sec. 2.1 we discuss the convective acceleration of vector fields. Section 2.2 contains properties of the functional $\|v_t + \nabla v v\|_{L^2}^2$ as a regularization term. The last part of Sec. 2 treats our choice of data term. Section 3 is dedicated to discussing our numerical minimization approach. Finally, we present experimental results in Section 4.

2 Model

2.1 Convective acceleration

Notation. Let $\Omega \subset \mathbb{R}^2$ be a bounded Lipschitz domain, $T > 0$, and set $E = (0, T) \times \Omega$. We denote points in Ω by $x = (x^1, x^2)^\top$. Let $\phi : E \rightarrow \Omega$ be a flow on Ω : That is, for fixed $t \in (0, T)$ the map $\phi(t, \cdot)$ is a diffeomorphism of Ω , while for fixed $x_0 \in \Omega$ the trajectory described by $\phi(\cdot, x_0)$ is smooth. The vector field on E that gathers the velocities of all trajectories associated to ϕ is defined by

$$u(t, x) = \phi_t(t, x_0) \Big|_{x_0 = \phi^{-1}(t, x)} \quad (3)$$

where ϕ_t is the partial derivative of ϕ with respect to its first argument and $\phi^{-1}(t, x)$ denotes the inverse of $\phi(t, \cdot)$ evaluated at x . Notice that $u(t, x)$ describes the velocity of the trajectory passing through x at the time t , so $u(s, x)$ in general corresponds to a different trajectory, if $s \neq t$.

Let $f : E \rightarrow \mathbb{R}^N$ be a possibly vector-valued quantity on E . We define the convective derivative of f along ϕ as

$$\begin{aligned} D_u f(t, x) &= \frac{d}{dt} f(t, \phi(t, x_0)) \Big|_{x_0 = \phi^{-1}(t, x)} \\ &= f_t(t, x) + \nabla f(t, x) u(t, x). \end{aligned} \quad (4)$$

Here $\nabla f = (f_{x^1}, f_{x^2}) \in \mathbb{R}^{N \times 2}$ is the spatial gradient of f . Notice that the convective derivative only depends on $u(t, x)$, thereby justifying the notation $D_u f$. Using the notation $\bar{\nabla} f = (f_t, \nabla f)$ and $\bar{u} = (1, u^\top)^\top$ we can write $D_u f = \bar{\nabla} f \bar{u}$. If we set $f = u$, the resulting vector field

$$D_u u = u_t + \nabla u u = \bar{\nabla} u \bar{u} \quad (5)$$

is called the *convective acceleration* of the flow ϕ .

Flows with vanishing convective acceleration. When proposing a new regularization term, it is always helpful to know when it is minimal. Therefore, we briefly discuss what conditions must be met so that a vector field u satisfies

$$D_u u = 0, \quad \text{for all } (t, x) \in E, \quad (6)$$

and give a few examples afterwards. Intuitively $D_u u = 0$ means that u satisfies a “flow constancy condition”. This is to be understood as an analogy to the widely known *brightness constancy condition*, which is frequently used to derive the optical flow equation. Note that with the notation introduced in the previous paragraph, for an image sequence f , the optical flow equation reads $D_u f = 0$.

From the definition of the convective derivative (4) and (3) it follows that

$$D_u u(t, \phi(t, x_0)) = \phi_{tt}(t, x_0),$$

for all $x_0 \in \Omega$ and $t \in (0, T)$. Now fix an $x_0 \in \Omega$ and, for simplicity, denote the trajectory $\phi(\cdot, x_0)$ by $\gamma(\cdot)$, so that we have $D_u u = \gamma''$. Assuming that γ' does not vanish on $(0, T)$, we can write

$$\begin{aligned} \gamma'' &= \frac{d}{dt} \left(|\gamma'| \frac{\gamma'}{|\gamma'|} \right) \\ &= \frac{d|\gamma'|}{dt} \frac{\gamma'}{|\gamma'|} + |\gamma'| \frac{d}{dt} \frac{\gamma'}{|\gamma'|}. \end{aligned}$$

Since the unit tangent vector $\gamma'/|\gamma'|$ is always perpendicular to its derivative, the convective acceleration γ'' is a sum of two mutually orthogonal vectors. Such a sum can only vanish, if both vectors vanish. Due to the assumption that $\gamma' \neq 0$, we conclude that

$$\begin{aligned} \frac{d|\gamma'|}{dt} &= 0, \quad \text{and} \\ \frac{d}{dt} \frac{\gamma'}{|\gamma'|} &= 0. \end{aligned}$$

This proves the quite intuitive fact that, if a flow has vanishing convective acceleration, then all its trajectories must be straight lines and have constant speed. Below we give a few examples of such vector fields.

Example 1. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function. The two vector fields $v_1, v_2 : E \rightarrow \mathbb{R}^2$ defined by

$$v_1(t, x^1, x^2) = \begin{pmatrix} 0 \\ g(x^1) \end{pmatrix}, \quad v_2(t, x^1, x^2) = \begin{pmatrix} g(x^2) \\ 0 \end{pmatrix}$$

do not depend on time and satisfy $\nabla v_i v_i = 0$. Thus, by formula (5) they have vanishing convective acceleration. Their integral curves are parallel to the x_2 - and x_1 -axes, respectively.

Example 2. Let $0 \notin \Omega$. Denote by e_r the unit vector in radial direction and let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function with period 2π . The vector field $v : E \rightarrow \mathbb{R}^2$ which in polar coordinates is given by

$$v(t, r, \varphi) = g(\varphi) e_r$$

has zero convective acceleration. In contrast to the previous example the integral lines of v are not all mutually parallel.

Example 3. The vector fields in the previous examples were all constant in time. Time-dependent examples can be constructed from time-dependent solutions u of the inviscid Burgers' equation in one space dimension ([6], Section 3.4). That is, let $u : E \rightarrow \mathbb{R}$ solve (6) with $E = (0, T) \times \mathbb{R}$. Then

$$v_1(t, x^1, x^2) = \begin{pmatrix} u(t, x^1) \\ 0 \end{pmatrix}, \quad v_2(t, x^1, x^2) = \begin{pmatrix} 0 \\ u(t, x^2) \end{pmatrix}$$

satisfy $D_{v_i} v_i = 0$ while having partial time derivatives that do not vanish identically. As in the first example the integral curves are mutually parallel straight lines. The main difference, however, is that two trajectories passing through the same point in space but at different times, might do so at different speeds.

Relation to curvature. Given a curve $\gamma : (0, T) \rightarrow \Omega$, we can define its curvature as the normal part of the arclength derivative of the unit tangent vector

$$\kappa_\gamma = \left(\frac{\gamma'}{|\gamma'|} \right)^\perp \cdot \frac{1}{|\gamma'|} \frac{d}{dt} \frac{\gamma'}{|\gamma'|} = \frac{\gamma'^\perp \cdot \gamma''}{|\gamma'|^3},$$

where $(a, b)^\perp = (-b, a)^\top$. If γ is arc length parametrized, that is $|\gamma'| \equiv 1$, then the expression simplifies to $\kappa_\gamma = |\gamma''|$, because in that case $\gamma'^\perp = \gamma''/|\gamma''|$.

For a given $x_0 \in \Omega$ we can apply the formula above to the trajectory $t \mapsto \phi(t, x_0)$ of a flow ϕ . Recall that in this case γ' becomes u and γ'' becomes $D_u u$. Thus we get

$$\kappa_{\phi(\cdot, x_0)} = \frac{u^\perp \cdot D_u u}{|u|^3}.$$

In particular, whenever $|u(t, \phi(t, x_0))| \equiv 1$ on some subinterval of $(0, T)$, then the norm of the convective acceleration $|D_u u|$ coincides with the absolute value of the curvature of $\phi(\cdot, x_0)$ on this subinterval

$$\kappa_{\phi(\cdot, x_0)} = \pm |D_u u|.$$

2.2 Convective regularization

In this section we study some properties of the functional

$$u \mapsto \frac{1}{2} \|D_u u\|_{L^2}^2 \tag{7}$$

as a regularization term for optical flow. By $\|\cdot\|_{L^2}$ we mean the norm of $L^2(E, \mathbb{R}^2)$. Below we continue to use this shorthand whenever convenient.

Interpretation of the convective term. Functional (7) has two interpretations, one as a smoothing term and another one as a projection term. Recalling the bar-notation we introduced in Section 2.1, the integrand of (7) can be written as a quadratic form in the partial derivatives of u

$$|D_u u|^2 = \bar{\nabla} u^1 \bar{u} \bar{u}^\top (\bar{\nabla} u^1)^\top + \bar{\nabla} u^2 \bar{u} \bar{u}^\top (\bar{\nabla} u^2)^\top. \tag{8}$$

Observe that the diffusion tensor

$$\bar{u} \bar{u}^\top = \begin{pmatrix} 1 & u^\top \\ u & u u^\top \end{pmatrix}$$

is a projection matrix onto the span of \bar{u} composed with a scaling by $|\bar{u}|^2$. Therefore, minimization of (7) leads to smoothing of the vector field u only in direction \bar{u} , where precedence is given to regions where the magnitude of u is relatively large. Clearly, since $\bar{u} \neq (1, 0, 0)^\top$ in general, the proposed convective regularization term is not a purely temporal regularizer, but it also enforces spatial smoothness of the vector field in a way that is consistent with the motion. Figure 2 illustrates this behavior.

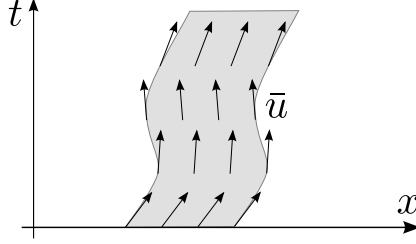


Figure 2: The shaded area represents the space-time signature of a moving object. The arrows not only indicate the space-time velocity \bar{u} of the object but also the direction of diffusion enforced by the tensor $\bar{u}\bar{u}^\top$. As opposed to the situation with the time derivative (Figure 1), object boundaries are respected.

On the other hand, the roles of \bar{u} and $\bar{\nabla}u^i$ can be exchanged so that (8) rewrites as a sum of two quadratic forms in \bar{u}

$$|D_u u|^2 = \bar{u}^\top (\bar{\nabla}u^1)^\top \bar{\nabla}u^1 \bar{u} + \bar{u}^\top (\bar{\nabla}u^2)^\top \bar{\nabla}u^2 \bar{u}. \quad (9)$$

The matrix $(\bar{\nabla}u^i)^\top \bar{\nabla}u^i$ projects onto the span of $\bar{\nabla}u^i$ and scales by the squared magnitude of $\bar{\nabla}u^i$. From this point of view minimizing (7) amounts to forcing \bar{u} to be orthogonal to the directions of greatest change of both u^1 and u^2 , where precedence is again given whenever this change is relatively drastic. Note that the outer product of a vector with itself is always a positive semidefinite matrix, so that there can be no cancellations in (8), (9).

The concurrence of the two different interpretations above is made explicit by looking at the Euler-Lagrange equation of (7), given by

$$((\nabla u^1)^\top \bar{\nabla}u^1 + (\nabla u^2)^\top \bar{\nabla}u^2) \bar{u} - \bar{\nabla} \cdot (\bar{\nabla}u \bar{u} \bar{u}^\top) = 0. \quad (10)$$

Here, the space-time divergence operator $\bar{\nabla} \cdot$ acts along rows.

Variational properties. In the following we adopt the notation $A \lesssim B$ whenever $A \leq B$ holds up to multiplication by a positive constant. According to the basic estimate

$$\|D_u u\|_{L^2}^2 \lesssim \|u_t\|_{L^2}^2 + \|u\|_{L^\infty}^2 \|\nabla u\|_{L^2}^2$$

a seemingly natural space over which to minimize, for a given image sequence f , the optical flow functional

$$\mathcal{F}(u) = \|D_u f\|_{L^2}^2 + \alpha \|D_u u\|_{L^2}^2$$

would be $X = H^1(E, \mathbb{R}^2) \cap L^\infty(E, \mathbb{R}^2)$. Now, in order to show existence of minimizers using the direct method [5], we need to ensure existence of a minimizing

sequence which converges. This is typically done by establishing a coercivity condition of the type

$$\mathcal{F}(u) \gtrsim \|u\|_X^p - b$$

for some $b \geq 0$ and $p \geq 1$. In the examples above we have seen, however, that there are vector fields u with vanishing convective acceleration but nonzero partial derivatives. This implies that not even the inequality

$$\|D_u u\|_{L^2}^2 \gtrsim \|\bar{\nabla} u\|_{L^2}^p - b$$

holds true. Therefore, in order to guarantee existence of minimizers we consider the functional

$$\begin{aligned} \mathcal{E} : H^1(E, \mathbb{R}^2) &\rightarrow [0, +\infty] \\ \mathcal{E}(u) &= \|\lambda D_u f\|_{L^2}^2 + \alpha \|D_u u\|_{L^2}^2 + \beta \|\bar{\nabla} u\|_{L^2}^2, \end{aligned} \quad (11)$$

where $\alpha, \beta > 0$ and $\lambda : E \rightarrow \mathbb{R}_+$ is a weighting function which is specified in Sec. 2.3. In addition we have to make the assumption introduced in [16] that the partial derivatives of f are linearly independent in $L^2(E)$, that is

$$\langle f_{x^1}, f_{x^2} \rangle_{L^2} < \|f_{x^1}\|_{L^2} \|f_{x^2}\|_{L^2}. \quad (12)$$

Proposition 1. *Let $f \in W^{1,\infty}(E)$ satisfy (12) and assume that $\lambda \in L^\infty(E)$ is such that $\text{ess inf } \lambda > 0$. Then the minimization of \mathcal{E} over $H^1(E, \mathbb{R}^2)$ has at least one solution.*

Proof. We show that \mathcal{E} is proper, coercive and lower semicontinuous with respect to the weak topology of $H^1(E, \mathbb{R}^2)$. That \mathcal{E} is proper follows from nonnegativity and the assumptions on f and λ .

Next we prove the coercivity estimate

$$\mathcal{E}(u) \gtrsim \|u\|_{H^1}^2 - b$$

for $b \geq 0$. We obviously have $\mathcal{E}(u) \geq \beta \|\bar{\nabla} u\|_{L^2}^2$. So it remains to show that $\mathcal{E}(u) \gtrsim \|u\|_{L^2}^2 - b$. Denoting the average of u by $u_E = 1/|E| \int_E u$ we have

$$\begin{aligned} \|u\|_{L^2}^2 &\lesssim \|u_E\|_{L^2}^2 + \|u - u_E\|_{L^2}^2 \\ &\lesssim \|u_E \cdot \nabla f\|_{L^2}^2 + \|\bar{\nabla} u\|_{L^2}^2 \\ &\lesssim \|u \cdot \nabla f\|_{L^2}^2 + \|(u - u_E) \cdot \nabla f\|_{L^2}^2 + \|\bar{\nabla} u\|_{L^2}^2 \\ &\lesssim \|D_u f\|_{L^2}^2 + \|f_t\|_{L^2}^2 + \|\bar{\nabla} u\|_{L^2}^2 \\ &\lesssim \mathcal{E}(u) + \|f_t\|_{L^2}^2, \end{aligned} \quad (13)$$

which proves coercivity. There are three main ingredients in the above estimate. The first one is a quadratic variant of the triangle inequality

$$\|v + w\|^2 \leq 2\|v\|^2 + 2\|w\|^2.$$

The second one

$$\|u_E\|_{L^2}^2 \lesssim \|u_E \cdot \nabla f\|_{L^2}^2$$

uses the assumption of linear independence of f_{x^1} and f_{x^2} and is proved in [16, p. 29]. The third ingredient is the Poincaré inequality $\|u - u_E\|_{L^2} \lesssim \|\bar{\nabla} u\|_{L^2}$

([9], Theorem 8.11). Also note that the fourth inequality in (13) requires the assumption $f \in W^{1,\infty}(E)$, while the last one uses $\text{ess inf } \lambda > 0$.

Let $\{u_n\} \subset H^1(E, \mathbb{R}^2)$ converge to \hat{u} in the weak topology of $H^1(E, \mathbb{R}^2)$. In particular, $\bar{\nabla} u_n$ converges weakly in $L^2(E, \mathbb{R}^6)$ to the corresponding gradient $\bar{\nabla} \hat{u}$. By the compact embedding of $H^1(E, \mathbb{R}^2)$ into $L^2(E, \mathbb{R}^2)$, up to choosing a subsequence we may assume that u_n also converges strongly in L^2 to \hat{u} . As is made apparent in (8), the expression $|D_u u|^2$ is convex in $\bar{\nabla} u$ for fixed u , and \mathcal{E} is always nonnegative. We can therefore apply a standard result ([5], Theorem 3.23) to obtain that \mathcal{E} is lower semicontinuous. \square

There is no reason to expect minimizers of \mathcal{E} to be unique. The following example demonstrates that \mathcal{E} is not convex in general.

Example 4. Let v_1, v_2 be the two vector fields from Example 1 with g being the identity, that is

$$v_1(t, x^1, x^2) = \begin{pmatrix} 0 \\ x^1 \end{pmatrix}, \quad v_2(t, x^1, x^2) = \begin{pmatrix} x^2 \\ 0 \end{pmatrix}$$

They satisfy $D_{v_1} v_1 = D_{v_2} v_2 \equiv 0$, while

$$D_w w = \frac{1}{4} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix},$$

where $w = (v_1 + v_2)/2$. We conclude

$$0 = \|D_{v_1} v_1\|_{L^2}^2 = \|D_{v_2} v_2\|_{L^2}^2 < \|D_w w\|_{L^2}^2 = \frac{T}{16} \int_{\Omega} |x|^2 dx$$

so that the functional $u \mapsto \|D_u u\|_{L^2}^2$ cannot be convex. This clearly implies that there are α, β, f such that \mathcal{E} is not convex.

2.3 Data term and contrast invariance.

Contrast invariance is a useful property of image processing operators. Let h be a change of contrast, that is, a differentiable real function with $h' > 0$. An operator A is called contrast invariant, if it commutes with all such contrast changes:

$$A \circ h(f) = h \circ A(f)$$

for all images f ([1], Sec. 2.3).

A similar property can be postulated for operators estimating image motion, since an order-preserving rearrangement of the grey values of an image sequence should certainly not change velocities. Therefore, we call an operator A that maps image sequences f to velocity fields *contrast invariant*, if it satisfies

$$A(f) = A \circ h(f).$$

However, a typical optical flow model of the form

$$f \mapsto \arg \min_u \{ \|D_u f\|_{L^p}^p + \alpha \mathcal{R}(u) \} \quad (14)$$

does not have this property. A simple counterexample is multiplication by a positive number $h(f) = cf$. For a model like (14) this change of contrast effectively amounts to dividing the regularization parameter α by c^p and therefore

definitely influences the result. There is, however, also a local effect of this contrast dependence, which is easy to confirm experimentally. Consider a scene with a dark background and two moving objects which are similar in shape and size and which move at similar velocities, but which have significantly different average brightnesses. This is precisely the case in the example of Figure 3. Then, the velocities of the darker object, i.e. the one with lower values of f , will be regularized more strongly than those of the brighter one, if a data term like the above is used.

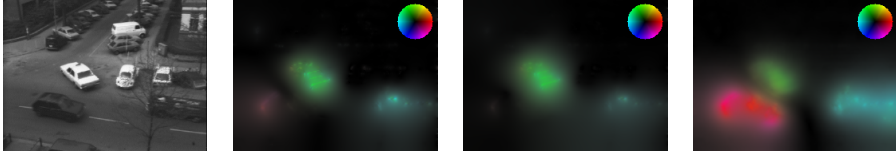


Figure 3: Effect of the data term weighting for \mathcal{E} with $\alpha = 0$. When it is omitted, the data term inappropriately favors high-contrast objects. Shown are, from left to right, one frame of f , the corresponding results without weighting and $\beta = 0.2$, without weighting and $\beta = 0.5$, and weighted with $\omega = \sqrt{|\bar{\nabla} f|^2 + \epsilon^2}$ and $\beta = 0.005$.

However, contrast dependence is not an inherent property of optical flow. Observe that, if u solves $D_u f = 0$, then it also solves $D_u(h \circ f) = 0$, because $D_u(h \circ f) = h' D_u f$ and $h' > 0$ by assumption. Therefore, it should rather be regarded a side-effect of the variational regularization approach. Contrast invariance can be restored by weighting the data term with $\lambda = 1/\omega$, where ω is a positively 1-homogeneous function that depends only on the first order derivatives of f . It then follows that

$$\frac{D_u(h \circ f)}{\omega(\bar{\nabla}(h \circ f))} = \frac{h' D_u f}{h' \omega(\bar{\nabla} f)} = \frac{D_u f}{\omega(\bar{\nabla} f)}. \quad (15)$$

Such weighted data terms have already been used, although it seems with a different reasoning. In [8, 20], for example, the weight was chosen to be $\omega = \sqrt{|\nabla f|^2 + \epsilon^2}$. In this paper we experiment with the weight $\omega = \sqrt{|\bar{\nabla} f|^2 + \epsilon^2}$. With this choice the data term essentially measures the orthogonal projection of \bar{u} onto the span of $\bar{\nabla} f$, since

$$\frac{D_u f}{\omega} \approx \frac{\bar{\nabla} f}{|\bar{\nabla} f|} \cdot \bar{u}.$$

The effectiveness of this choice is illustrated in Figure 3, in which the relative velocities of the vehicles are recovered correctly even though they have widely different contrast levels.

Finally, observe that weights which do not depend on first derivatives of f (only), like $\omega = |f| + \epsilon$ for instance, do not yield full contrast invariance. While they work fine for the special case of rescaling $h(f) = cf$, they fail in general since the factor h' in (15) does not cancel.

3 Numerical solution

We now seek to formulate a numerical approach for the minimization of the functional \mathcal{E} . As shown above, the convective regularization term is nonconvex and leads to nonlinear optimality conditions. Therefore, we propose an iterative scheme to arrive at an adequate flow field.

Iterative scheme. Consider the functional

$$\mathcal{G}(u, w) = \|\lambda D_u f\|_{L^2}^2 + \alpha \|D_w u\|_{L^2}^2 + \beta \|\bar{\nabla} u\|_{L^2}^2,$$

which satisfies $\mathcal{E}(u) = \mathcal{G}(u, u)$. For fixed $w \in L^\infty(E, \mathbb{R}^2)$, the mapping $u \mapsto \mathcal{G}(u, w)$ is a convex and quadratic functional, and $\mathcal{G}(u, w) < +\infty$ for any $u \in H^1(E)$. Our iterative method then reads as follows.

- Given an image sequence f and parameters $\beta_0, \alpha_1, \beta_1$:
- Find an initial guess u_0 by minimizing $\mathcal{G}(u, 0)$ over u , with $\alpha = 0, \beta = \beta_0$. This induces isotropic regularization in time and space, corresponding to a time-regularized variant of the Horn-Schunck model.
- Compute u_{k+1} by minimizing over the variable u the functional $\mathcal{G}(u, u_k)$ with parameters $\alpha = \alpha_1, \beta = \beta_1$.

Each of these steps corresponds to finding an optical flow field using anisotropic regularization with the diffusion tensor $\alpha \bar{w} \bar{w}^\top + \beta \text{Id}$, where Id is the 3×3 identity matrix. Therefore, this particular scheme is based on the diffusion interpretation of the convective regularization term, as reflected in (8). Indeed, upon fixing w the first term in the complete Euler-Lagrange equation (10) does not appear. Even though convergence of the iterative scheme is not guaranteed, in all of our experiments the solution u_k stabilized after only a few iterations. Typically by $k = 3$ or 4 the difference between successive iterates became negligible.

By the properties of \mathcal{G} listed above, each of these minimization problems can be performed by solving the corresponding linear Euler-Lagrange equation

$$\bar{\nabla} \cdot (\bar{\nabla} u (\alpha \bar{w} \bar{w}^\top + \beta \text{Id})) - \lambda^2 (\nabla f \cdot u) \nabla f^\top = \lambda^2 f_t \nabla f^\top, \quad (16)$$

coupled with natural boundary conditions

$$n \cdot (\bar{\nabla} u (\alpha \bar{w} \bar{w}^\top + \beta \text{Id})) = 0 \quad \text{on } \partial E,$$

where n is the outward normal to E . Our approach of discretizing (16) is specified next.

Discretization. The model has been discretized using a multilinear finite element formulation on a regular rectangular grid, so that each vertex in the grid corresponds to one pixel of one video frame. The grid used is such that the spacing between nodes in space is one, while the spacing in time (from frame to frame) may be smaller. We emphasize that the relation of the two grid step sizes implicitly controls the relative amount of regularization done in time and space, since the derivatives of the functional scale accordingly. Usually, using the same spacing will lead to results that are highly over-regularized in time. In

the presented examples, the time step was divided by a factor of 8 with respect to that of space.

Both the input f and the unknown vector field u are represented by their nodal values, leading to continuous piecewise multilinear functions. Since the Euler-Lagrange equation (16) is actually a system for the two components u^1, u^2 , the resulting linear system is represented by a block matrix, with the different blocks describing the interactions between components. The entries of these blocks are assembled from the integrals over different elements arising in the weak formulation, as per finite element practice [13]. Such integrals are computed exactly through an adequate Gauss quadrature. Note that boundary conditions do not need to be imposed explicitly.

A consequence of representing f by a piecewise multilinear function is that the corresponding partial derivatives f_{x^i}, f_t appearing in the right hand side $\lambda^2 f_t \nabla f^\top$ of (16) do not belong to the same finite element space as f . However, our finite element matrices act on the nodal values of functions in this space, so an interpolation step is needed to obtain functions with the appropriate smoothness. To this end, we have performed an L^2 projection of these partial derivatives into the space of piecewise multilinear functions, whose nodal values are then multiplied together and used as the right hand side for our linear system. Computationally each projection requires the solution of an additional linear system, which needs to be done only once.

4 Experiments

In Figures 4, 5 and 6 we present some examples of flows computed through the numerical scheme of the previous section, approximating minimizers of

$$\mathcal{E}(u) = \|\lambda D_u f\|_{L^2}^2 + \alpha \|D_u u\|_{L^2}^2 + \beta \|\bar{\nabla} u\|_{L^2}^2,$$

which are compared with corresponding results with basic isotropic regularization, that is, minimizers of

$$\mathcal{H}(u) = \|\lambda D_u f\|_{L^2}^2 + \beta \|\bar{\nabla} u\|_{L^2}^2.$$

In all these examples results were computed over an interval of 30 frames.

In the sequences used, the time derivative of the vector field at a given position is likely not appropriate, since isolated objects travel across the images. In the results with isotropic regularization, some of the disadvantages of basic quadratic regularizers are apparent. With lower regularization parameters (Figures 4(b) and 5(b)) inner edges of rigidly moving objects are visible. With higher regularization parameter values, the support of the vector field is enlarged and interactions between distinct objects moving ensue, creating cancellations when the objects move in different directions (Figure 4(c)) or artificial reinforcement when they move in similar directions (Figure 5(c)). However, with convective regularization both of these disadvantages can be avoided (Figures 4(d) and 5(d)).

Additionally, one needs to choose the regularization constant ϵ for the weight $\sqrt{|\bar{\nabla} f|^2 + \epsilon^2}$ in the data term. Clearly, choosing a too large ϵ will diminish the effect of the weighting scheme, while choosing it too small may overemphasize very low contrast regions and potentially amplify noise. Since the 8-bit input

images were normalized so that their values are of the form $k/255$ with $k = 0 \dots 255$, we chose $\epsilon = 0.01$. That is, ϵ is roughly of the size of the smallest nonzero derivative that may appear in the data.

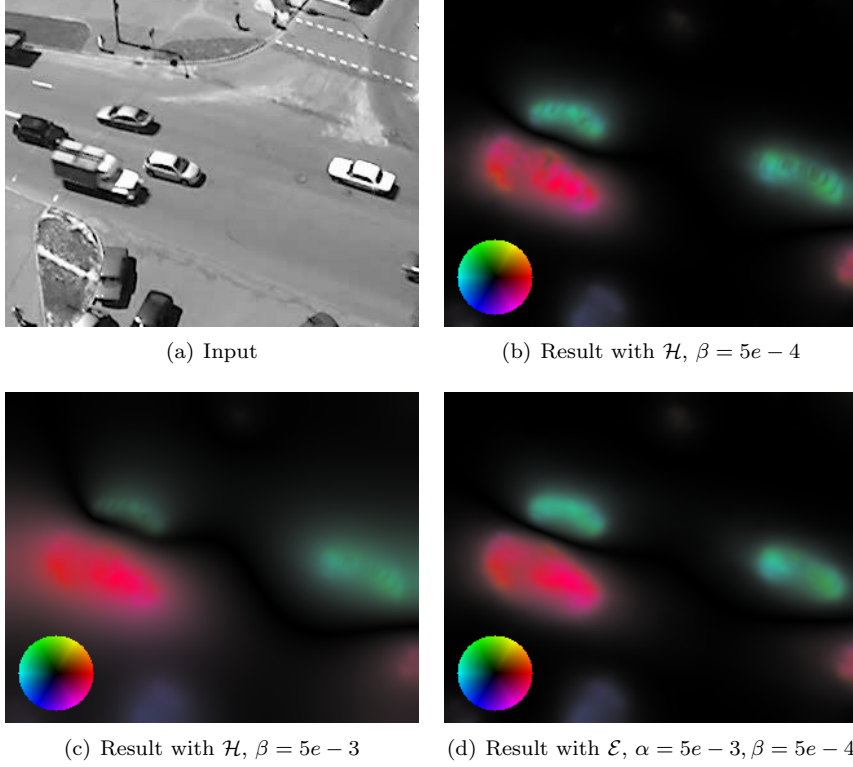


Figure 4: Comparison of isotropic and convective regularization.

In our experiments the parameter β_1 may be chosen quite small to preserve the boundaries of objects, while using a relatively large α_1 for the convective regularization term. By using the convective derivative, the information from other frames is used in a consistent way to correctly fill the flow field inside objects even in the absence of texture. The initial isotropic parameter β_0 for starting the iterative scheme was chosen to be $\beta_0 = \alpha_1$ in all cases.

5 Conclusion

In this article we presented an entirely continuous optical flow model that is based on the assumption that velocities should vary smoothly not at fixed points in space but along motion trajectories. The resulting regularization term has a variety of different interpretations. First, this term penalizes the proper acceleration of objects moving according to the velocity field. Second, it has a direct relation to the curvature of motion trajectories. Finally, when minimized it acts on the vector field not only as a projection term but also as an edge-preserving anisotropic regularizer.

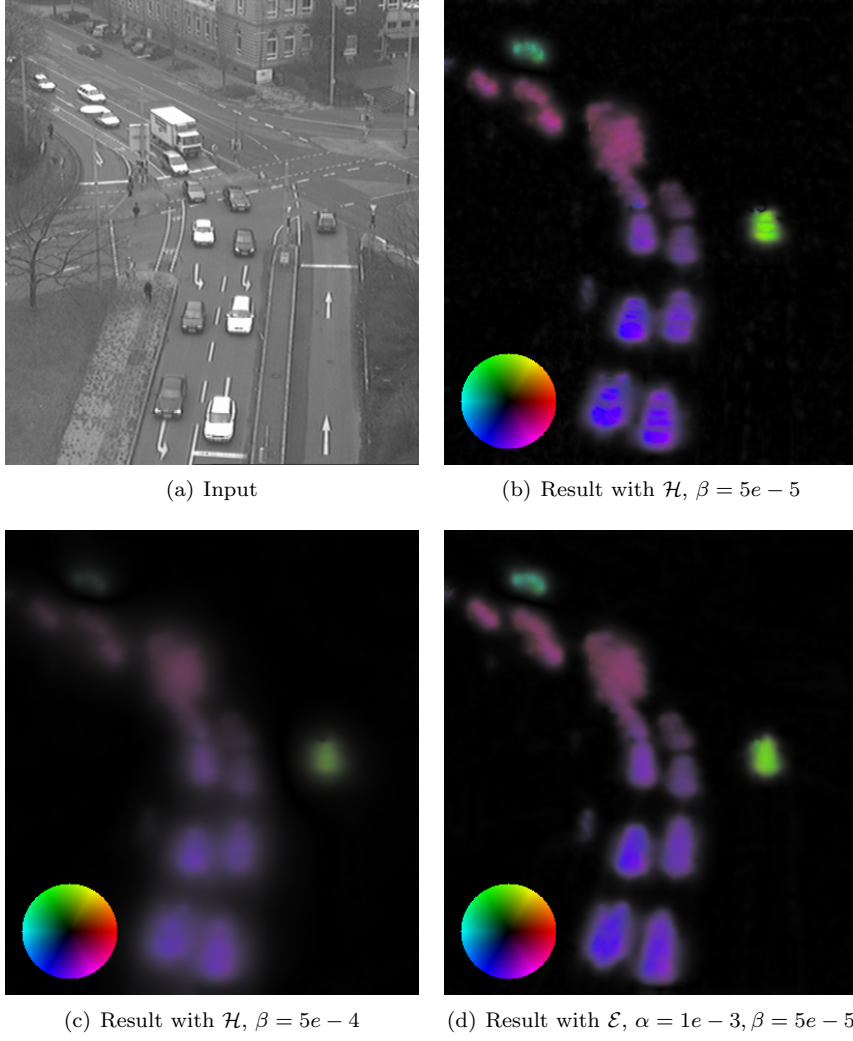


Figure 5: Comparison of isotropic and convective regularization.

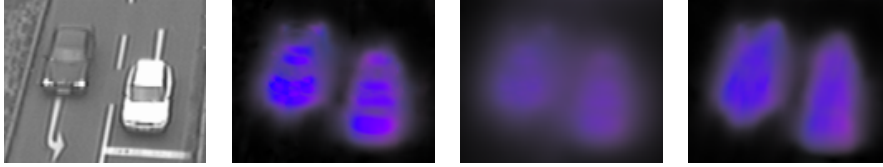


Figure 6: Cutout from Figure 5. The use of convective regularization allows to obtain nearly-constant optical flow field inside the objects without excessively enlarging the support.

We tested the convective regularization term by incorporating it into a reference optical flow model and comparing the results with and without the new term. In order to make the changes as apparent as possible and to avoid potential side effects we tried to keep the complexity of the reference model at a minimum. Therefore, the resulting functional (11) should mainly be viewed as a prototype of an optical flow model with convective regularization which can be improved upon by future research.

Acknowledgements. This work has been supported by the Austrian Science Fund (FWF) within the national research network Geometry + Simulation (project S11704, Variational Methods for Imaging on Manifolds). Additionally, JI acknowledges support from the FWF Doctoral Program Dissipation and Dispersion in Nonlinear PDEs (W1245). We wish to thank Enrique C. Vargas for inspiring exchanges. Our finite element implementation was carried out in the `Quocmesh` library,¹ AG Rumpf, Institute for Numerical Simulation, Universität Bonn. For our experiments, we have used footage from the datasets of Lund University² (sequence “Conflicts-20100628-1344, camera 1”) and Universität Karlsruhe³ (sequences “taxi” and “dt-passat-part2”).

References

- [1] L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations of image processing. *Arch. Ration. Mech. Anal.*, 123(3):199–257, 1993.
- [2] F. Amat, E. W. Myers, and P. J. Keller. Fast and robust optical flow for time-lapse microscopy using super-voxels. *Bioinformatics*, 29(3):373–380, 2013.
- [3] K. Chaudhury and R. Mehrotra. A trajectory-based computational model for optical flow estimation. *IEEE Trans. on Robot. Autom.*, 11(5):733–741, Oct 1995.
- [4] J.P. Cocquerez, L. Chanas, and J. Blanc-Talon. Simultaneous inpainting and motion estimation of highly degraded video-sequences. In J. Bigun and T. Gustavsson, editors, *Image Analysis*, volume 2749 of *Lecture Notes in Computer Science*, pages 685–692. Springer Berlin Heidelberg, 2003.
- [5] B. Dacorogna. *Direct methods in the calculus of variations*, volume 78 of *Applied Mathematical Sciences*. Springer, New York, 2 edition, 2008.
- [6] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, second edition, 2010.
- [7] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

¹<http://numod.ins.uni-bonn.de/software/quocmesh/>

²http://www.tft.lth.se/video/co_operation/data_exchange/

³ftp://ftp.ira.uka.de/pub/vid-text/image_sequences/

- [8] S.-H. Lai and B.C. Vemuri. Reliable and efficient computation of optical flow. *Int. J. Comput. Vision*, 29(2):87–105, 1998.
- [9] E. H. Lieb and M. Loss. *Analysis*, volume 14 of *Graduate Studies in Mathematics*. American Mathematical Society, second edition, 2001.
- [10] H.-H. Nagel. Constraints for the estimation of displacement vector fields from image sequences. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, IJCAI 83, pages 945–951, 1983.
- [11] H.-H. Nagel. Extending the ‘oriented smoothness constraint’ into the temporal domain and the estimation of derivatives of optical flow. In *Proceedings of the First European Conference on Computer Vision*, ECCV 90, pages 139–148, New York, NY, USA, 1990. Springer-Verlag New York, Inc.
- [12] T. Preusser, M. Droske, C.S. Garbe, A. Telea, and M. Rumpf. A phase field method for joint denoising, edge detection, and motion estimation in image sequence processing. *SIAM J. Appl. Math.*, 68(3):599–618, 2008.
- [13] A. Quarteroni and A. Valli. *Numerical Approximation of Partial Differential Equations*, volume 23 of *Springer Series in Computational Mathematics*. Springer Verlag, Berlin, 2008.
- [14] A. Salgado and J. Sánchez. Temporal constraints in large optical flow estimation. In R. Moreno-Díaz, F. Pichler, and A. Quesada-Arencibia, editors, *Computer Aided Systems Theory – EUROCAST 2007*, volume 4739 of *Lecture Notes in Computer Science*, pages 709–716. Springer Berlin Heidelberg, 2007.
- [15] B. Schmid, G. Shah, N. Scherf, M. Weber, K. Thierbach, C. Campos Pérez, I. Roeder, P. Aanstad, and J. Huisken. High-speed panoramic light-sheet microscopy reveals global endodermal cell dynamics. *Nat. Commun.*, 4:2207, 2013.
- [16] Ch. Schnörr. Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *Int. J. Comput. Vision*, 6:25–38, 1991.
- [17] S. Volz, A. Bruhn, L. Valgaerts, and H. Zimmer. Modeling temporal coherence for optical flow. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1116–1123, Nov 2011.
- [18] J. Weickert and Ch. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *Int. J. Comput. Vision*, 45(3):245–264, 2001.
- [19] J. Weickert and Ch. Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *J. Math. Imaging Vision*, 14:245–255, 2001.
- [20] H. Zimmer, A. Bruhn, and J. Weickert. Optic flow in harmony. *Int. J. Comput. Vision*, 93(3):368–388, 2011.